Astronomy & Astrophysics manuscript no. paper_arxiv
January 21, 2014
©ESO 2014

A new approach to multi-frequency synthesis in radio interferometry

H. Junklewitz[1,2], M. R. Bell[1], and T. Enßlin[1,2]

[1] Max-Planck Institut für Astrophysik (Karl-Schwarzschild-Str. 1, D-85748 Garching, Germany)
[2] Ludwig-Maximilians-Universität München (Geschwister-Scholl-Platz 1, D-80539 München, Germany)

January 21, 2014

arXiv:1401.4711v1 [astro-ph.IM] 19 Jan 2014

ABSTRACT

We present a new approach to multi-frequency synthesis in radio astronomy. Using Bayesian inference techniques, the new technique estimates the sky brightness and the spectral index simultaneously. In principle, the bandwidth of a wide-band observation can be fully exploited for sensitivity and resolution, currently only limited by higher order effects like spectral curvature. Employing this new approach, we further present a multi-frequency extension to the imaging algorithm RESOLVE. In simulations, this new algorithm outperforms current multi-frequency imaging techniques like MS-MF-CLEAN.

1. Introduction

In radio astronomy, multi-frequency observations are widely used for many different purposes. Examples include the investigation of spectral line emission, analyzing continuous synchrotron spectra of radio sources, the determination of Faraday rotation in polarization imaging, correlating brightness structures at largely different wavelengths or improving the sensitivity and coverage of an interferometer without introducing more antennas into the array (for a review see e.g. Taylor et al. 1999). Most of these only become possible with observations that span over many frequencies and data analysis methods that can handle this kind of observations.

For the remainder of this work, we focus entirely on multi-wavelength radio continuum studies in total intensity. Neither spectral line observations, polarization imaging, nor wavelengths studies beyond pure radio observations will be a direct topic. In Sec. 4, we will comment on possible extensions of the presented work into other domains of multi-frequency radio astronomy.

Historically in radio interferometry, the term *multi-frequency synthesis* is mainly used to denote techniques that focus on the direct combination of single-instrument observations at different frequencies to improve the resolution and sensitivity of a radio interferometer (see Conway et al. (1990) and references therein). For this purpose, a number of methods have been devised, most notably double deconvolution (Conway et al. 1990) and multi-frequency CLEAN (Sault & Wieringa 1994). These methods usually work by Taylor-expanding a spectral model function around a reference frequency. We will comment on these methods in more detail in Sec. 2.1.

Conversely, investigating the spectral behavior of a radio synchrotron source on its own usually is achieved through totally independent standard imaging of the surface brightness in the sky at a number of frequencies. Resolution is kept uniform over all frequencies to correctly recombine the images at individual frequencies. Subsequently, the spectral parameters are determined by fitting a function through all the single frequency images, usually assuming a power-law shaped spectral evolution.

Only recently, implementations of MF-CLEAN also make it possible to constrain the spectral properties (see Reid & CASA Team 2010), effectively starting to merge both multi-frequency high resolution imaging and spectral analysis into one algorithm. It is this combined notion, how we like to understand and use the term multi-frequency synthesis.

Further development of multi-frequency imaging techniques is paramount for being able to fully exploit the data from the new generation of radio telescopes, such as the upgraded VLA, LOFAR, the SKA pathfinder missions or ultimately the SKA itself (see e.g. Garrett 2012). Their unprecedented, broadband frequency coverages of many GHz of possible bandwidths and previously unknown frequency regimes offer many advances in astrophysical and cosmological sciences. But at the same time, they are a challenge for current data analysis methods (see Sec. 2.1). One part of the challenge rises from the fact that current algorithms might not meet the expected sensitivity or fidelity because previously acceptable models and approximations break down with the quality and quantity of data now available. The other side is that the huge quantities in which new data sets come, represent a huge computational challenge on their own, regardless of the algorithms used. It should be emphasized that this study focuses entirely on the first part of the challenge (for this topic, see also Sec. 4).

In this paper, we present a new approach that uses Bayesian statistical inference techniques to combine a *simultaneous* statistical estimation of the sky brightness and the spectral behavior with the benefits in resolution and sensitivity of classical multi-frequency synthesis. Conceptually, this new approach has a number of advantages over standard methods. In principle, the full bandwidth of the observation is used for maximum theoretical sensitivity. No subsequent imaging at all single frequencies is employed and, thereby, less reconstruction artifacts are introduced and no artificial downgrading of resolution is necessary.

Our model for the spectral behavior is not approximated through Taylor-expansions around a frequency, and in the moment only is limited by possible higher order effects like spectral curvature. In this way, our approach is conceptually very similar to the new method of *Faraday Synthesis* (Bell & Enßlin 2012) in polarization imaging (see Sec. 2.2 for more details). Furthermore, the spatial correlation of the parameters describing the spectral behavior is used to constrain and improve the estimation of spectral properties, and can further be viewed as a new scientific result on its own. Finally, it is possible to approximate the statistical uncertainty of the spectral index estimate in addition to that of the total intensity.

To demonstrate the viability of our new approach, we present a multi-frequency extension to the radio extended emission imager RESOLVE (Junklewitz et al. 2013) as an alternative to multi-scale-multi-frequency CLEAN (Rau & Cornwell 2011), the standard method for wide-band multi-frequency synthesis of extended radio sources.

For our derivations, we will often refer to the work presented in Junklewitz et al. (2013), henceforth Res2013.

2. Theory & Algorithm

In this section, we first briefly review the current status of multi-frequency synthesis techniques (Sec. 2.1), then develop the new approach (Sec. 2.2) based on *information field theory* (Enßlin et al. 2009) and the inference framework presented in Res2013. We further present a multi-frequency-extension to the imaging algorithm RESOLVE (see Res2013) using the developed framework (Sec. 2.2).

2.1. Multi-frequency aperture synthesis and spectral index reconstruction

Aperture synthesis is the technique of connecting an array of telescopes in such a way that we can effectively synthesize a combined instrument with a much larger aperture and therefore resolution (Ryle & Hewish 1960; Thompson et al. 1986). It can be shown that for observations at a single frequency ν_0, and under the assumption of measuring the sky as flat in a plane tangent to the phase center of the observation, such a radio interferometer approximatively takes incomplete samples of the Fourier transformed brightness distribution in the sky (Thompson et al. 1986):

$$V(u,v,\nu_0) \approx W(u,v,\nu_0) \int \mathrm{d}l \; \mathrm{d}m \; I_0(l,m,\nu_0) \; \mathrm{e}^{-2\pi i(ul+vm)}. \tag{1}$$

For our purposes, working under this assumption suffices for the analysis undertaken in this paper. For a detailed derivation of (1) consider (Thompson et al. 1986). The quantity $V(u,v)$ is called the *visibility* function. The coordinates u and v are vector components describing the distance between a pair of antennas in an interferometric array, where this distance is usually referred to as a *baseline*. They are given in numbers of wavelengths, with u and v usually parallel to geographic east-west and north-south, respectively. The coordinates l and m are a measure of the angular distance from the phase center along axes parallel to u and v, respectively. $W(u,v)$ is a sampling function defined by the layout of the interferometric array. It is zero throughout most of the u,v-space, apart from where measurements have been made where it is taken to be unity.

The visibility function is what our instrument measures, but we are actually interested in the brightness distribution of the source in the sky. Unfortunately, an inversion of (1) gives us not the true brightness distribution, but its convolution with the inverse Fourier transform of the sampling function, better known as the *dirty beam* $I_{db} = \mathcal{F}^{-1}W$:

$$I_\mathrm{D} = \mathcal{F}^{-1}V = \mathcal{F}^{-1}W\mathcal{F}I = I_{db} * I. \tag{2}$$

Here, we have introduced a symbolic Fourier operator $\mathcal{F}_{kx} = \exp(-i(ul+vm))$ with $x = (l,m)$ and $k = (u,v)$, the common notation I_D, *dirty image*, for the simple Fourier inversion of the visibilities, and the symbol $*$ to denote a convolution operation.

For imaging at a single frequency, one usually proceeds using a deconvolution algorithm aiming to solve (2) approximatively. The most common algorithm is CLEAN (Högbom 1974), or one of its many variants (Clark 1980; Schwab 1984; Sault & Wieringa 1994; Cornwell 2008; Rau & Cornwell 2011), which basically model the sky brightness to be a collection of delta peak point sources that are iteratively assigned by searching for the peak values in the dirty image. But alternatives exist, especially for imaging extended emission, like the Maximum Entropy Method (Cornwell & Evans 1985), Adaptive Scale Pixel decomposition (Bhatnagar & Cornwell 2004) or the recently published Bayesian extended emission imager RESOLVE (Res2013), which will be used in this paper in an expanded multi-frequency version (see Secs. 2.2 and 3).

Originally, radio telescopes observed with relatively narrow bandwidths and at only a few different frequencies. The standard approach for imaging the spectral properties of a source since then is to take single frequency observations, or observations averaged over close channels for higher sensitivity, image them separately, and then fit a spectral model to the observations. Since continuous synchrotron emission is known to show a power-law spectrum (Rybicki & Lightman 1985), for many astrophysical purposes such a model is sufficient:

$$I(l,m,\nu) = I_0(l,m,\nu_0) \; \left(\frac{\nu}{\nu_0}\right)^{-\alpha}. \tag{3}$$

Sometimes, higher order spectral deviations are modeled with a second term in the exponential of (3):

$$I(l,m,\nu) = I_0(l,m,\nu_0) \; \left(\frac{\nu}{\nu_0}\right)^{(-\alpha+\log\nu/\nu_0 \;\beta)}, \tag{4}$$

effectively enhancing the linear function in $\log I$ vs. $\log\nu$ space (3) to a quadratic polynomial. This model is referred to as spectral curvature.

Simultaneous multi-frequency imaging at different frequencies has been introduced to radio astronomy upon the realization that observations at many frequencies, if stacked together appropriately, can improve the sampling in the uv-plane (see Conway et al. (1990) and references therein). That way, the sensitivity of a radio interferometric observation can be enhanced considerably. Because the uv-coordinates are measured in numbers of wavelengths, the

same interferometer samples different parts of the Fourier space at different frequencies. An interferometer includes $N(N-1)/2$ baselines and therefore uv points, where N is the number of antennas. In theory, observing at a number of frequencies N_f enhances the measured baselines to $N_f N(N-1)/2$, which is the equivalent effect of introducing roughly $\sqrt{N_f}$ extra antennas (Conway et al. 1990).

This approach was first developed by Conway et al. (1990) for mid-sized fractional bandwidths of around $\pm 10\%$, together with the method of double deconvolution to mitigate spectral errors when using CLEAN for a spectrally combined data set. Later, double deconvolution was developed into the more advanced multi-frequency CLEAN (MF-CLEAN) (Sault & Wieringa 1994) and multi-frequency multi-scale CLEAN (MF-MS-CLEAN) (Rau & Cornwell 2011).

All these methods assume the spectral dependence to be a power-law, like Eqs. (3) or (4) and propose to approximate it during the CLEAN-like-deconvolution process with a Taylor-expansion. To our knowledge, current implementations use a few terms of a Taylor-expansion in ν or $\log \nu$ around a reference frequency ν_0. In Rau & Cornwell (2011) it is discussed that a direct decomposition into one or two terms of a polynomial in $\log I$ vs. $\log \nu$ would actually be the most accurate representation[1], but for current implementations this is discarded by the authors because of numerical instabilities.

During the CLEAN deconvolution process, the iteratively updated sky model is used to also update the coefficients of the spectral Taylor expansion. In this way, spectral index or even spectral curvature can be constrained by these coefficients. Care must be taken, since this expansion is usually stopped after a few terms and might not be valid over large bandwidths. Because of this, the only implementation of MF-MS CLEAN known to the authors to date, within the radio astronomical software package CASA (Reid & CASA Team 2010), is considered to be experimental and on a shared-risk basis with regard to the spectral reconstructions with a higher number of terms.

For modern wide-band data sets from the new generation of instruments, spanning several GHz of bandwidths, this means that probably even more emphasis must be put onto the handling of the spectral effects, either by invoking higher order terms in the expansions (see Sault & Wieringa 1994; Rau & Cornwell 2011) or by shifting to a different approach, where the sky brightness and its spectral properties are fully considered simultaneously, and estimated to fit the entire data using some global minimization function. This last approach was actually mentioned by Conway et al. (1990), but found to be unnecessarily complicated and computationally expensive for the typically modest fractional bandwidths at the time.

We will now consider the latter approach and present a statistical solution.

2.2. A multi-frequency extension to the algorithm RESOLVE

In the course of this section, we often refer to the detailed derivations layed out in Sec. 2 and App. 1 of Res2013 that form the basis from which we derive our multi-frequency algorithm.

We start by summarizing (1) and (3) into a multi-frequency measurement model:

$$d(k,\nu) = W(k,\nu) \int dx\ \mathcal{F}(k,x) I(x,\nu) + n(k,\nu) \tag{5}$$

where again $x = (l,m)$ and $k = (u,v)$, the term $n(k,\nu)$ introduces measurement noise, the data $d(k,\nu)$ has been introduced, which is basically the visibility function with measurement noise, and the spectral dependence of the source $I(x,\nu)$ is kept general for the moment.

In order to simplify notation and to identify (5) as a general inverse inference problem as analyzed in Res2013, we henceforth drop all explicit dependence on k and x and combine all known instrumental effects into a response function $R_\nu := W(k,\nu)\mathcal{F}$, leading to

$$d_\nu = R_\nu I_\nu + n_\nu. \tag{6}$$

There are many instrumental effects beyond the sampling in the uv-plane given by $W(k,\nu)$ like an antenna sensitivity pattern or an direction dependent, variable sampling. For this work, we stay with the basic definition $R_\nu := W(k,\nu)\mathcal{F}$ and refer the reader to Res2013 concerning the possibility of including other effects within this framework. We just emphasize that, without loss of generality, most of these effects can be in principle included[2]. We also assume the instrument to be fully calibrated, and thus the response R_ν to be known.

In accordance with standard radio interferometric literature (Thompson et al. 1986), we assume Gaussian noise statistics, mainly induced by the antenna electronics and independent between measurements at different frequencies and time steps of the observation (Thompson et al. 1986). Henceforth, the noise n_ν will be assumed to be drawn from a multivariate, zero mean Gaussian distribution of dimension n_d:

$$\begin{aligned}\mathcal{P}(n) &= \mathcal{G}(n,N) \\ &:= \frac{1}{\det(2\pi N_\nu)^{1/2}}\ \exp\left(-\frac{1}{2} n_\nu^\dagger N_\nu^{-1} n_\nu\right).\end{aligned} \tag{7}$$

Solving (6) exactly for I_ν is not possible, since all information is lost on the Fourier modes not sampled by R_ν. This is just a different way of stating that a direct Fourier inversion of (1) yields the dirty image and not an exact representation of the sky brightness.

Instead, the aim is to find a statistical estimate for the *most probable* sky brightness signal, given all observational and noise constraints. In Res2013, it was shown in detail how this can be done in a Bayesian statistical framework for a radio astronomical data model like (1). We briefly repeat the main points, and otherwise refer the reader to Res2013.

To find an optimal statistical estimate for the sky brightness signal I_ν, we regard it as a random field with certain

[1] Actually, for the spectral model (4), this would be no approximation at all since it simply is a quadratic polynomial in $\log I$ vs. $\log \nu$ space.

[2] In principle this approach can be extended into a full RIME (Radio Interferometer Measurement Equation), considering all Stokes parameters and instrumental gains (Smirnov 2011a,b).

a priori statistical properties expressed in the prior distribution $P(I)$, but fully constrained by the data through the statistics of the *likelihood* distribution $P(d|I)$. The likelihood distribution summarizes how the data are obtained with a measurement of the true sky brightness signal, and for our problem can be expressed as

$$\mathcal{P}(d|I) = \mathcal{G}(d_\nu - R_\nu I_\nu, N_\nu), \tag{8}$$

which is a Gaussian over $d_\nu - R_\nu I_\nu$ with the covariance structure of the uncorrelated noise $N_{kk'\nu} = \delta_{kk'}\sigma^2_{k\nu}$.

Prior and likelihood statistics can be combined into the *posterior* distribution $P(I|d) \propto P(d|I)P(I)$ that holds the important information of how much the sky brightness signal is statistically constrained by the data. From there, an estimate for the signal can be obtained by calculating a suitable statistic of the posterior, most prominently its mean or its mode, corresponding to the minimization of different error norm measures between the signal and its estimate (see Res2013; or Jaynes (2003), Caticha (2008) and Lemm (1999), Enßlin et al. (2009) for a comprehensive review of Bayesian statistics or inference on fields respectively).

The exact choice of an appropriate inference algorithm at this stage largely depends on the complexity of the problem (i.e. the posterior). For the problem at hand, since the likelihood is already known (8), this comes down to the question of the prior statistics of I_ν.

The most general way conceivable would be *not* to explicitly model the spectral dependence of $I(x,\nu)$ at all, as for instance in (3) or (4). Instead, $I(x,\nu)$ can be interpreted as a three dimensional continuous field and should be inferred as a whole from the entire data. In general, let us assume not only point-like but also extended emission in the sky, and some kind of extended structure in spectral space as well. In such a setting, $I(x,\nu)$ could be set *a priori* as a statistical field with an unknown and probably non-isotropic cross-correlation structure in the combined sky and spectral space. Such a complex, unknown cross-correlation structure at the one hand complicates the problem enormously, but could also be used to guide the reconstruction, if correctly estimated with the field itself.

A number of statistical methods have already been developed to handle simpler problems of signal reconstructions with unknown but isotropic correlation structure, many of them solving the problem for Gaussian fields using information field theory (Enßlin & Frommert 2011; Enßlin & Weig 2010) or Gibbs-Sampling Monte-Carlo methods (Sutter et al. 2012; Karakci et al. 2013), or recently also for log-normal fields (Enßlin & Frommert 2011; Enßlin & Weig 2010; Oppermann et al. 2013; Greiner 2013; Selig & Enßlin 2013). Most notably this method was also used to create RESOLVE (Res2013). A full combined spatial and spectral reconstrcution as outlined above, would require substantial further development and is outside of the scope of this work.

We can choose a more direct strategy, still residing within our approach of statistical inference, but fix a spectral model and infer instead the spectral index (or curvature) as a field on its own. For the rest of the paper, we will choose the model (3) for the simplicity of the approach, and for a functional similarity with the algorithm RESOLVE that makes it very natural to include into a combined method.

RESOLVE works under the assumption that the extended surface brightness at a single frequency is *a priori* assumed as a random field drawn from log-normal statistics (Res2013). For our multi-frequency problem, this basically turns (6) into

$$\begin{aligned} d_\nu &= R_\nu I_\nu + n_\nu \\ &= R_\nu \left[I_0(l,m,\nu_0) \left(\frac{\nu}{\nu_0}\right)^{-\alpha} \right] + n_\nu \\ &= R_\nu \left[\rho_0 \, \mathrm{e}^{s(l,m)} \left(\frac{\nu}{\nu_0}\right)^{-\alpha} \right] + n_\nu, \end{aligned} \tag{9}$$

where s is a Gaussian random field (such that the logarithm of e^s is a Gaussian random field again), and ρ_0 is a constant to e.g. normalize the system to the right units. Although, the frequency dependence of s was not explicitly shown in Res2013, the derivation of RESOLVE implicitly assumed the algorithm to work for a single frequency in the way presented here. RESOLVE assumes the spatial correlation of an extended source in the sky to be reflected by the covariance of the Gaussian random field s, which is unknown *a priori*, and thus estimated from the data itself together with the sky brightness. The covariance S of a Gaussian field is equivalent to its two-point correlation function $S(l,m) = \langle s(l)s(m)^\dagger \rangle$ and is handled as such by RESOLVE in form of the power spectrum $S(k,k') = \langle s(k)s(k')^\dagger \rangle$, which is the Fourier transformation of the correlation function[3]. A deeper analysis of this can be found in Res2013.

We now make the central assumption that the spectral index α can be modeled *a priori* as a Gaussian random field with its own spatial correlation structure in the sky. At least for an extended source, we have every reason to assume that the spectral index should be a field with spatial extension itself. Observational constraints strongly imply that typical extended radio structures show as well extended and smooth spectral index structures, for instance radio halos and relics of galaxy clusters (Feretti et al. 2012), radio galaxy lobes (Kassim et al. 2005), or supernova remnants (Green 2001).

In Res2013, we argued extensively that a Gaussian random field would be the ideal choice for a signal prior of an extended field with a priori unknown correlation structure, as long as the field is not assumed to vary strongly on orders of magnitude and not necessarily needs to be positive definite. Both constraints apply very well to known spectral index maps, where variations usually do not reach even one order of magnitude, and nothing prevents the spectral index in principle to change the sign. For details of these arguments see Res2013, we now proceed under this assumption.

If we rewrite (3) only slightly,

$$\begin{aligned} I(l,m,\nu) &= I_0(l,m,\nu_0) \left(\frac{\nu}{\nu_0}\right)^{-\alpha} \\ &= I_0(l,m,\nu_0) \, \mathrm{e}^{-\ln(\nu/\nu_0)\alpha}, \end{aligned} \tag{10}$$

[3] We actually assume the spatial correlation to be *a priori* rotationally and translationally invariant, and thus the power spectrum to be diagonal $S(k,k') = \langle s(k)s(k')^\dagger \rangle = (2\pi)^{n_s}\delta(k-k')P_s(|k|)$. However, this does not imply that the correlation structure must be invariant under any transformation *a posteriori* as well. A more detailed discussion can be found in Res2013.

it reveals that, if α has a Gaussian prior, (10) naturally turns into a model for a log-normal prior, only different in shape (and more complicated) from (9) because of the term $-\ln(\nu/\nu_0)$.

It is important to note that we have not specified $I_0(l,m,\nu_0)$ yet. Thus, at this point, our inference approach to multi-frequency synthesis is in principle compatible with any method that reconstructs and deconvolves the surface brightness $I_0(l,m,\nu_0)$ at a single reference frequency ν. At least as long as it seems consistent with the source of interest to assume that the spectral index is an extended and spatially correlated field. In an extreme case, like single, unresolved point sources, this method probably will not yield optimal results.

For this paper, we take the choice to combine (9) and (10) into one single method, where we assume our double log-normal measurement model to be

$$d_\nu = R_\nu \left[\rho_0 \, \mathrm{e}^{s(l,m) - \ln(\nu/\nu_0)\alpha(l,m)}\right] + n_\nu, \tag{11}$$

with ρ_0 again a constant, from now on set to one w.l.o.g., and the signal fields s and α having Gaussian prior distributions $P(s)$ and $P(\alpha)$:

$$\begin{aligned}\mathcal{P}(s) &= \mathcal{G}(s,S) \\ &= \frac{1}{\det(2\pi S)^{1/2}} \, \mathrm{e}^{-\frac{1}{2} \, s^\dagger S^{-1} s}, \end{aligned} \tag{12}$$

$$\begin{aligned}\mathcal{P}(\alpha) &= \mathcal{G}(\alpha,A) \\ &= \frac{1}{\det(2\pi A)^{1/2}} \, \mathrm{e}^{-\frac{1}{2} \, \alpha^\dagger A^{-1} \alpha}. \end{aligned} \tag{13}$$

We now write down a posterior distribution for each signal field, while the other field (and its covariance) is assumed to be known, held constant and regarded as part of two distinct versions of the response operator R_ν, henceforth called $R_{(\mathrm{s})} = W(k,\nu)\mathcal{F}\left(\mathrm{e}^{-\ln(\nu/\nu_0)\alpha} \circ\right)$ and $R_{(\alpha)} = W(k,\nu)\mathcal{F}\left(\mathrm{e}^{s(l,m)} \circ\right)$ (where the symbol $\circ$ denotes where the field needs to be inserted that the operator acts on):

$$\mathcal{P}(s|d) \propto \mathcal{G}(d - R_{(\mathrm{s})}\mathrm{e}^s, N) \; \mathcal{G}(s,S), \tag{14}$$

$$\mathcal{P}(\alpha|d) \propto \mathcal{G}(d - R_{(\alpha)}\mathrm{e}^{-\ln(\nu/\nu_0)\alpha}, N) \; \mathcal{G}(\alpha,A). \tag{15}$$

As in Res2013, it is not possible to calculate the posterior mean for either of the two signal fields s or α without invoking complicated or expensive perturbation or sampling methods (see Res2013). We therefore continue to use the procedure already presented there, and calculate the posterior maximum to estimate both signals

$$\begin{aligned} m_\mathrm{s} &= \mathrm{argmax}_s \mathcal{P}(s|d), \\ m_\alpha &= \mathrm{argmax}_\alpha \mathcal{P}(\alpha|d). \end{aligned} \tag{16}$$

In signal inference, this procedure is called Maximum A Posteriori (MAP) (Jaynes 2003). The resulting fix-point equations from Eqs. (16) need to be solved numerically using a non-linear optimization scheme. For this, we resort to the same implementations as in Res2013 (see App. 2 therein).

With this choice, we basically extend RESOLVE to a multi-frequency algorithm by integrating a second complete RESOLVE step for the spectral index into the method, and iterating between the statistical estimation of s with its covariance S, and α with its covariance A. As outlined above, we always hold one of the fields (and their respective covariances) as constant and regard them as part of the response during this process.

It should be emphasized again that the RESOLVE step for the spectral index could in principle be combined with any other method to reconstruct $I_0(l,m,\nu_0)$ at a single frequency ν_0.

The exact equations that need to be solved to calculate (16) for either of the two fields and estimate their power spectra (i.e. their correlation structure in form of their Gaussian covariances, see above) are layed out in App. A and derived rigorously in Res2013.

3. Tests

We have integrated multi-frequency capability into the existing implementation of RESOLVE and tested the algorithm using simulated data[4]. The code is written in PYTHON using the signal inference library NIFTY (Selig et al. 2013), for details of the implementation we refer the reader to App. 2 of Res2013.

As in Res2013, we constructed simulated observations with the tool MAKEMS[5] using a realistic uv-coverage from a VLA observation in its A-Configuration. The VLA samples the uv-plane non-uniformly at irregular intervals, and the response includes thereby a convolutional gridding and degridding operator using a Kaiser-Bessel kernel (for details see App. 2 in Res2013). We simulated observations over a range of 2 GHz, with 20 separate frequency channels. The observations are short snapshots of approximatively 20 minutes per frequency, with a total of 42 120 visibility measurements at each frequency channel (see Fig. 1). This setting leads to an especially sparse sampling of the uv-plane at a single frequency, but to a much better coverage for the combined multi-frequency data (see Fig. 1).

For all tests, the signal s is drawn from a Gaussian distribution, finally entering the formalism as an exact log-normal field[6]

In order to use a spectral index signal with some correlation to the brightness signal, we model the spatial source dependence of the spectral index in an ad hoc fashion.

For the spectral index signal α, we use a sum of two Gaussian fields (see Sec. 3.1), one being independently drawn, while the second is the surface brightness signal s itself, down-weighted with a suitable factor. This is done in order to introduce a spectral index signal with some correlation to the surface brightness signal, at least in an ad-hoc fashion. Typically, extended sources in radio astrophysics show a spatial cross-correlation between both, which is a

[4] To get access to the preliminary code prior to its envisaged public release, please contact henrikju@mpa-garching.mpg.de or ensslin@mpa-garching.mpg.de.

[5] See http://www.lofar.org/wiki/lib/exe/fetch.php?media=software:makems.pdf.

[6] In Res2013 it was demonstrated that RESOLVE also works beyond that on realistic signals drawn from real CLEAN maps.

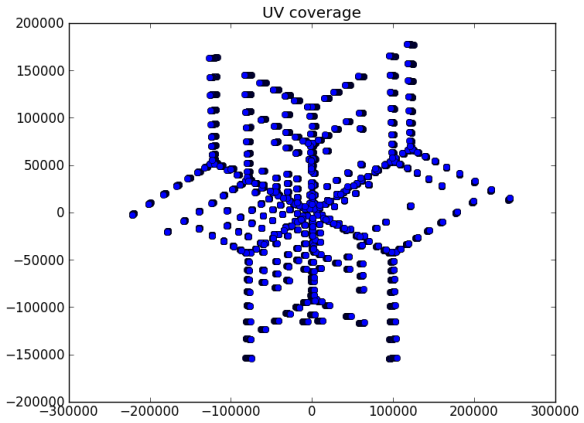

(a) Single frequency uv-coverage in units of # of wavelengths.

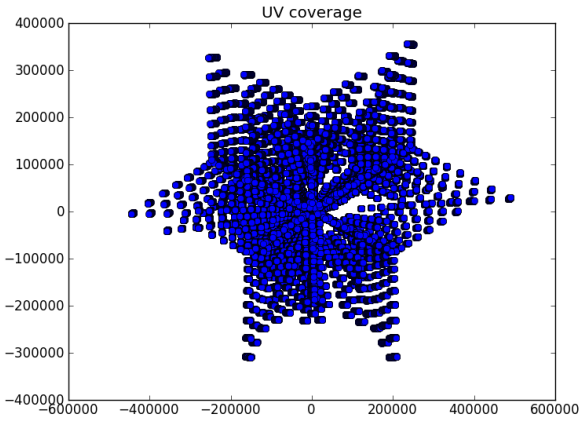

(b) Full multi-frequency uv-coverage in units of # of wavelengths.

Fig. 1. Single-frequency and full-frequency uv-coverage from all frequencies of a simulated 20 minutes snapshot observation in VLA-a configuration.

result of different physical processes underlying the structure and formation of these sources. A notable example are radio halos and relics in galaxy clusters (Feretti et al. 2012). Of course, our model is only an ad-hoc approximation for the proof of concept undertaken in this work. We also emphasize that such a cross-correlation between α and s is *not* exploited by RESOLVE (see Sec. 2.2 for an outlook).

The complex, Gaussian input noise variance in uv-space is equal for all visibilities and frequencies. As with single-frequency RESOLVE, the algorithm does not require equal noise variances and can in principle handle varying variances. The noise variance was set to a low[7] value of $\sigma^2 = 10^{-3}\mathrm{Jy}^2$.

In continuation with Res2013, we use a relative $\mathcal{L}_2$ - norm measure of the difference in the signal to the reconstruction to measure the accuracy of the estimate of both brightness and spectral reconstructions:

$$\delta_s = \sqrt{\left(\frac{\sum \left(\mathrm{e}^s - \mathrm{e}^m\right)^2}{\sum \left(\mathrm{e}^s\right)^2}\right)}, \tag{17}$$

$$\delta_\alpha = \sqrt{\left(\frac{\sum \left(\alpha - m_\alpha\right)^2}{\sum \left(\alpha\right)^2}\right)}, \tag{18}$$

where the sums are taken over all pixels of the reconstruction. For the motivation behind this choice see Sec. 3 in Res2013.

In Sec. 3.1 we focus exclusively on the reconstruction of the two signals s and α. Then, in Sec. 3.2, we compare the results from multi-frequency RESOLVE with standard imaging procedures. The reconstruction of the signal power spectra is discussed separately in Sec. 3.3.

3.1. Main test results

We start by showing the reconstruction of the presented simulated observations using multi-frequency RESOLVE. In Fig. A.1, a surface brightness and a spectral index signal are shown, together with the respective reconstructions obtained with RESOLVE and absolute difference maps of both signals to their reconstructions. We show the spectral index maps in full, and over-layed with a mask that focuses on the part of the observed field that contains the brightest part of the surface brightness signal. Later, in Fig. A.2, we show the reconstruction for different masks (along with a comparison to other methods, see Sec. 3.2). We choose the three different masks mainly qualitatively by visual comparison to show only the parts of the sky with a surface brightness signal above a certain threshold. The actual threshold values for the surface brightness were 2 Jy/px, 4 Jy/px, and 6 Jy/px. In Fig. A.1, the second conservative mask is used. The error measures are $\delta_I = 0.13$ for the surface brightness, and $\delta_\alpha = 0.35$ for the spectral index.

It can be seen that RESOLVE recovers very accurately the original surface brightness. For this particular choice of relative low noise, the algorithm succeeds in reconstructing even small scale features of the signal. The effects of the instrumental point spread function are successfully deconvolved. These findings are in perfect agreement with the results presented on single-frequency RESOLVE in Res2013.

The general structure of the spectral index is well reconstructed, even for outer regions where the brightness signal is very weak. Overall, small scale features are much better recovered in the inner regions, where the main brightness sources are located, as can be seen by comparison between the full and the masked images in Fig. A.1. It is expected that the quality of the spectral index reconstruction depends on the strength of the observed surface brightness, as is illustrated by the poor performance of standard methods in recovering any structures outside the strong source regions, presented in the next section (see Fig. A.2). In most real applications, the outer parts of the observed fields should therefore usually be not a focus of the investigation. We thus choose the relatively conservative mask in Fig. A.1 to highlight this important part of the reconstruction. Nevertheless, it should be noted that – at least for this low noise example with relatively high sensitivity due to the broad bandwidth – RESOLVE is able to reliably extrapolate

[7] In comparison to the signal strength in Fourier space, the chosen value ensures a high signal to noise ratio. The unit Jy was used here for convenience. Effectively, it stands for whatever units the simulated signal is interpreted to be given in.

its estimation also into the weaker regions around the main sources.

In addition, with RESOLVE, an uncertainty of the spectral index reconstruction can be estimated. As for the single frequency reconstructions presented in Res2013, the second derivative of the posterior is used to approximate its covariance D_α (for details, see App. A and Res2013). This way, the full estimate for the spectral index signal becomes

$$\alpha_x \approx (m_\alpha)_x \pm \left(\sqrt{D_\alpha}\right)_{xx} . \quad (19)$$

For the given set of simulated data, roughly 60 % of the original signal values within the unmasked regions in Fig. A.1 lie within a $1-\sigma$ interval , but only roughly 20 % lie within a $1-\sigma$ interval for the full spectral index image. Due to the non-linear nature of the inference problem (16), it is expected that the uncertainty estimate is not exactly what would be expected from a pure Gaussian covariance (i.e. 68 % in the $1-\sigma$ interval), the MAP estimate is not guaranteed to lie very close to the real posterior covariance (see Res2013 for this problem). It is no surprise that the estimate worsens for all the outer regions with only very weak brightness structures present, which explains the poor performance on the whole spectral index map. Furthermore, as discussed in Res2013, the calculation of an uncertainty estimate is computationally costly. Due to a lack of accessible computer power only for testing purposes, we stopped the calculations at some point and smoothed the outcome. This should not be a great problem, since the uncertainty is expected to be smooth and we are mainly interested in a proof of concept at this point.

3.2. Comparison to standard methods

In this section, we compare the results of multi-frequency RESOLVE with two standard methods: A straight forward power-law fit for (3) using single-frequency CLEAN images, and a MF-MS-CLEAN reconstruction for the surface brightness and the spectral index. The reconstructions were performed on the same simulated observation as in Sec. 3.1. All results were obtained using the radio-astronomical software package CASA (Reid & CASA Team 2010). The MF-MS-CLEAN reconstructions were obtained using 1500 iterations, a small gain factor of 0.1, uniform weighting and ten different multi-scales ranging from a single pixel to moderately large structures. We further used two terms for the Taylor expansion in ν (see Sec. 2.1) and only half the available bandwidth, since a larger frequency range seems to lie outside the convergence radius of the Taylor-expansion (see Sec. 2.1 used in the implementation of CASA.

In Fig. A.2, a comparison is shown between RESOLVE, a power-law fit, and MF-CLEAN results for the spectral index and differently strong masks. The full field is not shown, because neither the power-law fit nor the CLEAN reconstruction could recover any structures in the remaining regions. The error measures are listed in Tab. 1. It can be seen that for this example RESOLVE overall outperforms the other methods. The advantage of RESOLVE is more pronounced in the outer regions, where the surface brightness signal is weak, but RESOLVE still gives the best result even for the most central parts of the reconstruction (see also the full reconstruction in Fig. A.1).

Algorithm \Mask	liberal mask	medium mask	conservative mask
RESOLVE	0.34	0.33	0.32
power-law fit	132.79	0.94	0.48
MS-MF CLEAN	3.07	1.89	1.84

Table 1. $\mathcal{L}_2$ error measures for RESOLVE, the power law fit, and MS-MF CLEAN for the simulation and the reconstruction observation of Sec. 3.1 and the three different masks defined in Sec. 3.1.

Fig. A.3 shows a comparison of multi-frequency surface brightness reconstructions with RESOLVE and MS-MF-CLEAN for the simulated observation of Sec. 3.1. For this simulation, RESOLVE is more successful in reconstructing the overall structure, and especially recovers more of the small scales. It should be noted that this particular CLEAN image was achieved using the full bandwidth and coverage (other than for the spectral index images, as stated at the beginning of the section).

3.3. Power spectrum reconstructions

The power spectrum of the spectral index is reconstructed together with the field itself. It represents the spatial correlation of the spectral index over the observed sky. As explained in Sec. 2.2, we expect the typical spectral index structure of an extended (i.e. spatially correlated) radio source to be spatially correlated to itself. For all the details on power spectrum reconstructions we refer the reader to Res2013, since almost all derivations for standard single-frequency RESOLVE power spectrum reconstructions are valid as well for the spectral index.

In this section we only discuss the spectral index power spectrum, since the reconstructions for the brightness power spectra are identical to the ones already presented in Res2013. We simply note that, in principle, a multi-frequency reconstruction recovers structure on smaller scales, and thus, it is expected that multi-frequency RESOLVE should be able to map spatial power spectra up to higher Fourier modes (i.e smaller correlation structures).

A typical result from multi-frequency RESOLVE is illustrated in Fig. 2. It shows the original spectral index power spectrum of the Gaussian signal field, used in the simulated observations of Sec. 3.1, and its reconstruction that belongs to the same iteration step as the presented signal reconstructions earlier. The power spectrum is reconstructed relatively well, no power is lost on high modes (i.e. small scales), which is just a consequence of the fact that the simulated observation was conducted with low noise.

As for the brightness reconstructions in Res2013, we emphasize that the power spectrum should not only be viewed as a by-product of the algorithm in order to accurately estimate the spectral index signal. For instance, some astronomical objects show very distinct spectral index structures and can even be classified after this criterion. A prominent example might be radio halos and relics of galaxy clusters, both of which typically show steep spectral indices that evolve spatially roughly like the source itself (Feretti et al. 2012). Measuring the spatial power spectra of these objects might lead to a more exact and quantitative classification scheme. Another application might lie in the investigation of different physical processes within a single source that lead to spectrally very different regions. Esti-

mating the spectral correlation structure over such regions offers a new way of quantitative analysis of the interplay of these processes.

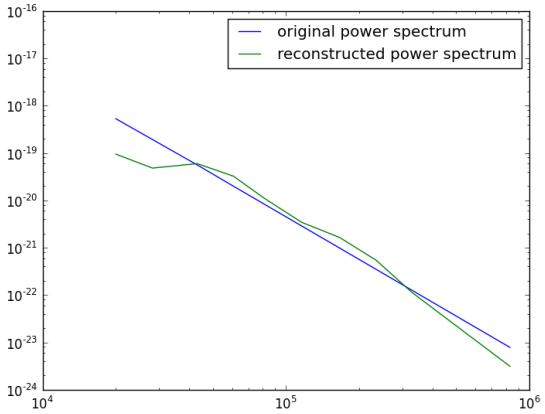

(a) Spectral power spectrum reconstruction

Fig. 2. Power spectrum reconstruction of the spectral index α for the reconstruction shown in Sec. 3.1 using multi-frequency RESOLVE.

4. Conclusion

We presented a multi-frequency extension to the imaging algorithm RESOLVE (Junklewitz et al. 2013). The combined algorithm is optimal for multi-frequency imaging of extended radio sources. It simultaneously estimates the surface brightness at a reference frequency, and the spectral index of the source. Within the assumption of a spectral index model, no further expansions or parameter-dependent modeling is used in the reconstruction. Multi-frequency RESOLVE is thus capable of exploiting the full bandwidth of a modern radio observation for maximum sensitivity and resolution, only limited by higher order spectral effects like spectral curvature.

Multi-frequency RESOLVE has been tested successfully using simulated observations of the VLA in its A-configuration. For the presented tests, the algorithm can outperform standard imaging methods in both surface brightness and spectral index estimations.

The algorithm uses a Gaussian prior and an effective log-normal model for the spectral index. This approach is not necessarily restricted to a combination with single-frequency RESOLVE, and can in principle be combined with any other imaging or deconvolution method for the surface brightness reconstruction.

For the sake of feasibility, many details of radio interferometric observations have been left out of the analysis. The response might realistically contain a number of additional effects, for instance an instrumental primary beam or wide-field and direction dependent effects. We refer the reader to Junklewitz et al. (2013), where many of these problems already have been discussed in the outlook.

In its current form, the presented algorithm has relatively high computational costs and numerical demands (for an analysis of algorithmic efficiency, see Junklewitz et al. 2013). In general, this is true for most Bayesian statistical inference algorithms. This could pose an obstacle for day-to-day applicability, since especially modern broad band data sets tend to be very large, and therefore are already a numerically challenge on their own. Since this study was centered on fundamental algorithmic development, numerical efficiency was not a focus. Future work might be needed to obtain a more efficient implementation of the algorithm.

For the future, it seems to be desirable to refrain completely from an explicit spectral model, and try to infer a full, non-parametric, three dimensional spectral intensity $I(l, m, \nu)$. Possibly, such a development could benefit greatly from reconstructing a full cross-correlation structure between the sky and spectral space. We leave this more complete approach for a future publication.

References

Bell, M. R. & Enßlin, T. A. 2012, A&A, 540, A80
Bhatnagar, S. & Cornwell, T. J. 2004, A&A, 426, 747
Caticha, A. 2008, ArXiv e-prints
Clark, B. G. 1980, A&A, 89, 377
Conway, J. E., Cornwell, T. J., & Wilkinson, P. N. 1990, MNRAS, 246, 490
Cornwell, T. J. 2008, IEEE Journal of Selected Topics in Signal Processing, 2, 793
Cornwell, T. J. & Evans, K. F. 1985, A&A, 143, 77
Enßlin, T. A. & Frommert, M. 2011, Phys. Rev. D, 83, 105014
Enßlin, T. A., Frommert, M., & Kitaura, F. S. 2009, Phys. Rev. D, 80, 105005
Enßlin, T. A. & Weig, C. 2010, Phys. Rev. E, 82, 051112
Feretti, L., Giovannini, G., Govoni, F., & Murgia, M. 2012, A&A Rev., 20, 54
Garrett, M. A. 2012, in From Antikythera to the Square Kilometre Array: Lessons from the Ancients
Green, D. A. 2001, in American Institute of Physics Conference Series, Vol. 558, American Institute of Physics Conference Series, ed. F. A. Aharonian & H. J. Völk, 59–70
Greiner, M. 2013, The Galactic Free Electron Density: A Bayesian Reconstruction, Master Thesis
Högbom, J. A. 1974, A&AS, 15, 417
Jaynes, E. T. 2003, Probability Theory: The Logic of Science
Junklewitz, H., Bell, M. R., Selig, M., & Enßlin, T. A. 2013, ArXiv e-prints (Res2013)
Karakci, A., Sutter, P. M., Zhang, L., et al. 2013, ApJS, 204, 10
Kassim, N. E., Lazio, T. J. W., Lane, W. M., et al. 2005, in X-Ray and Radio Connections, ed. L. O. Sjouwerman & K. K. Dyer
Lemm, J. C. 1999, ArXiv Physics e-prints
Oppermann, N., Selig, M., Bell, M. R., & Enßlin, T. A. 2013, Phys. Rev. E, 87, 032136
Rau, U. & Cornwell, T. J. 2011, A&A, 532, A71
Reid, R. I. & CASA Team. 2010, in Bulletin of the American Astronomical Society, Vol. 42, American Astronomical Society Meeting Abstracts #215, 479.04
Rybicki, G. B. & Lightman, A. P. 1985, Radiative processes in astrophysics.
Ryle, M. & Hewish, A. 1960, MNRAS, 120, 220
Sault, R. J. & Wieringa, M. H. 1994, A&AS, 108, 585
Schwab, F. R. 1984, AJ, 89, 1076
Selig, M., Bell, M. R., Junklewitz, H., et al. 2013, A&A, 554, A26
Selig, M. & Enßlin, T. 2013, ArXiv e-prints
Smirnov, O. M. 2011a, A&A, 527, A106
Smirnov, O. M. 2011b, A&A, 527, A107
Sutter, P. M., Wandelt, B. D., & Malu, S. S. 2012, ApJS, 202, 9
Taylor, G. B., Carilli, C. L., & Perley, R. A., eds. 1999, Astronomical Society of the Pacific Conference Series, Vol. 180, Synthesis Imaging in Radio Astronomy II
Thompson, A. R., Moran, J. M., & Swenson, G. W. 1986, Interferometry and synthesis in radio astronomy

Appendix A: Details of the algorithm

In this appendix, we repeat some of the basic derivations from Res2013, and derive the details of how reconstructing the spectral index using RESOLVE differs from the standard procedure presented there. In principle, many of the original derivations are still valid for multi-frequency RESOLVE, estimating s and α. The derivations for power spectrum reconstructions and uncertainty calculations stay especially close to the equations already presented in Res2013, and we do not repeat them explicitly here.

Appendix A.1: Reconstruction of the sky brightness signal field s

The estimation of the sky brightness at a reference frequency $I_0(l,m,\nu_0) = \rho_0 \mathrm{e}^{s(l,m)}$ can mostly be conducted with the standard single-frequency RESOLVE.

The only difference is the more complex response operator $R_{(\mathrm{s})} = W(k,\nu)\mathcal{F}\left(\mathrm{e}^{-\ln(\nu/\nu_0)\alpha}\circ\right)$. Dropping the explicit signal index for a moment and writing the specific measurement model (11) for the signal s with explicit operator and field indices

$$d_{k\nu} = R_{k\nu x}\ (\rho_0\ \mathrm{e}^{s})_x + n_{k\nu}, \tag{A.1}$$

reveals that the operator $R_{k\nu x}$ actually spans the single-frequency Fourier transform over all observed frequencies into a large, many-frequency data space[8]. Conversely, the adjoint response $R^\dagger$ now includes a sum over all frequencies to collapse everything back into the two dimensional sky at reference frequency. Since both operations are used in RESOLVE, it is by this procedure that the multi-frequency RESOLVE effectively uses uv-information at all frequencies to constrain the estimation of s. Of course, the quality of this reconstruction depends on the accuracy of the current estimation for α, used in the reconstruction step for s to define $R_{(\mathrm{s})}$.

Appendix A.2: Reconstruction of the spectral index α

The log-normal model for the spectral index, $\mathrm{e}^{-\ln(\nu/\nu_0)\alpha}$, only differs in the term $b = -\ln(\nu/\nu_0)$ from the model used for the sky brightness signal. This slightly complicates the calculations for the signal estimation with MAP and the power spectrum reconstruction and eventual uncertainty calculation. For both, derivatives of the posterior (15) with respect to α are needed (see Res2013), and this will add extra b-terms into the equations. The results of this are summarized in Sec. A.3.

As for the sky brightness signal s (see App. A.1), the response operator for the α reconstructions $R_{(\alpha)} = W(k,\nu)\mathcal{F}\left(\mathrm{e}^{s(l,m)}\circ\right)$ is more complex. Effectively, the b-term in the basic log-normal model $\mathrm{e}^{-\ln(\nu/\nu_0)\alpha}$ would prevent us from just defining our signal space to be the two-dimensional sky since $R_{(\alpha)}$ acts on $\mathrm{e}^{-\ln(\nu/\nu_0)\alpha}$. In the actual implementation of the multi-frequency RESOLVE algorithm, we circumvent this problem by assuming that $R_{(\alpha)}$ acts on α and all other terms, including the exponential operation of the log-normal model, are part of the operator[9]. In this way, $R_{(\alpha)}$ and $R^\dagger_{(\alpha)}$ can be understood to act in the same way as their respective counterparts for s do (i.e. they also span up and collapse into the full frequency data space).

Appendix A.3: Combined algorithm

Using our findings in Res2013 and in the previous subsections, multi-frequency RESOLVE comes down to solving iteratively two only slightly different, subsequent sets of equations:

Estimation of s

$$S_p^{-1}\ m_\mathrm{s} + \mathrm{e}^{m_\mathrm{s}}\cdot M_\mathrm{s}\mathrm{e}^{m_\mathrm{s}} - j_\mathrm{s}\cdot\mathrm{e}^{m_\mathrm{s}} = 0 \tag{A.2}$$

$$\begin{aligned}(D_\mathrm{s})_{xy} = {} & S_{p\ xy}^{-1} + \mathrm{e}^{(m_\mathrm{s})_x}\ (M_\mathrm{s})_{xy}\ \mathrm{e}^{(m_\mathrm{s})_y} \\ & + \mathrm{e}^{(m_\mathrm{s})_y}\int dz\ M_\mathrm{s}(x,z)\ \mathrm{e}^{(m_\mathrm{s})_z} \\ & - (j_\mathrm{s})_x\cdot\mathrm{e}^{(m_\mathrm{s})_x}\ \delta_{xy}\end{aligned} \tag{A.3}$$

$$p_i^s = \frac{\left(q_i^s + \frac{1}{2}\mathrm{tr}\left((m_\mathrm{s}{m_\mathrm{s}}^\dagger + D_\mathrm{s})S^{(i)}\right)\right)}{\left(\alpha_i^{(\mathrm{s-pr})} - 1 + \frac{\varrho_i}{2} + (Tp)_i\right)} \tag{A.4}$$

Estimation of α

$$A_p^{-1}m_\alpha + b\ \mathrm{e}^{bm_\alpha}\cdot M_\alpha\mathrm{e}^{bm_\alpha} - j_\alpha\cdot b\ \mathrm{e}^{bm_\alpha} = 0 \tag{A.5}$$

$$\begin{aligned}(D_\alpha)_{xy} = {} & A_{p\ xy}^{-1} + b^2\ \mathrm{e}^{b(m_\alpha)_x}\ (M_\alpha)_{xy}\ \mathrm{e}^{b(m_\alpha)_y} \\ & + b^2\ \mathrm{e}^{b(m_\alpha)_x}\int dz\ M_\alpha(x,z)\ \mathrm{e}^{b(m_\alpha)_z} \\ & - (j_\alpha)_x\cdot\mathrm{e}^{(m_\alpha)_x}\ \delta_{xy}\end{aligned} \tag{A.6}$$

$$p_i^\alpha = \frac{\left(q_i^\alpha + \frac{1}{2}\mathrm{tr}\left((m_\alpha{m_\alpha}^\dagger + D_\alpha)A^{(i)}\right)\right)}{\left(\alpha_i^{(\alpha-\mathrm{pr})} - 1 + \frac{\varrho_i}{2} + (Tp)_i\right)} \tag{A.7}$$

A rigorous derivation for all equations can be found in Res2013. The two sets of equations only differ in form by the $b = -\ln(\nu/\nu_0)$ - terms that show up in the spectral index reconstruction because of the derivatives used in order to calculate the MAP estimate. The quantities j_s, M_s and j_α, M_α are defined as

$$j_\mathrm{s} = R_\mathrm{s}^\dagger N^{-1} d, \tag{A.8}$$
$$M_\mathrm{s} = R_\mathrm{s}^\dagger N^{-1} R_\mathrm{s}, \tag{A.9}$$
$$j_\alpha = R_\alpha^\dagger N^{-1} d, \tag{A.10}$$
$$M_\alpha = R_\alpha^\dagger N^{-1} R_\alpha. \tag{A.11}$$

$S^{(i)}$, or $A^{(i)}$ are projection operators onto a band of Fourier modes denoted by the index i, while p_i are parameters to model the unknown power spectrum into a number of such

[8] As a reminder: In (A.1) we use an implicit convention to sum or taking the integral over repeated discrete or continuous indices, see Res2013 for details.

[9] This actually renders $R_{(\alpha)}$ a non-linear operator.

bands $S = \sum_i p_i S^{(i)}$, or $A = \sum_i p_i A^{(i)}$ (see Res2013 for details). The quantities q, $\alpha^{(\mathrm{pr})}$ and ϱ are parameters of a power spectrum prior for the signal or the spectral index, and T is an operator, which enforces a smooth solution of the power spectrum p_i.

Eqs. (A.2) and (A.5) are the fix point equations that need to be solved numerically to find a MAP signal estimate m_s or m_α for the current iteration. The second equations (A.3) and (A.6) result from calculating the second derivative of the respective posteriors for the signal estimates m_s or m_α, their inverses serve as an approximation to the signal uncertainty $D_\mathrm{s} = \langle (s - m_s)(s - m_s)^\dagger \rangle$ or $D_\alpha = \langle (\alpha - m_\alpha)(\alpha - m_\alpha)^\dagger \rangle$ at each iteration step. The last equations (A.4) and (A.7) represent an estimate for the signal power spectra (and therefore their autocorrelation functions), using the signal uncertainties D_s or D_α to correct for missing signal power in the current estimates m_s or m_α . The iteration is stopped after a suitable convergence criterion is met (see App. 2 in Res2013).

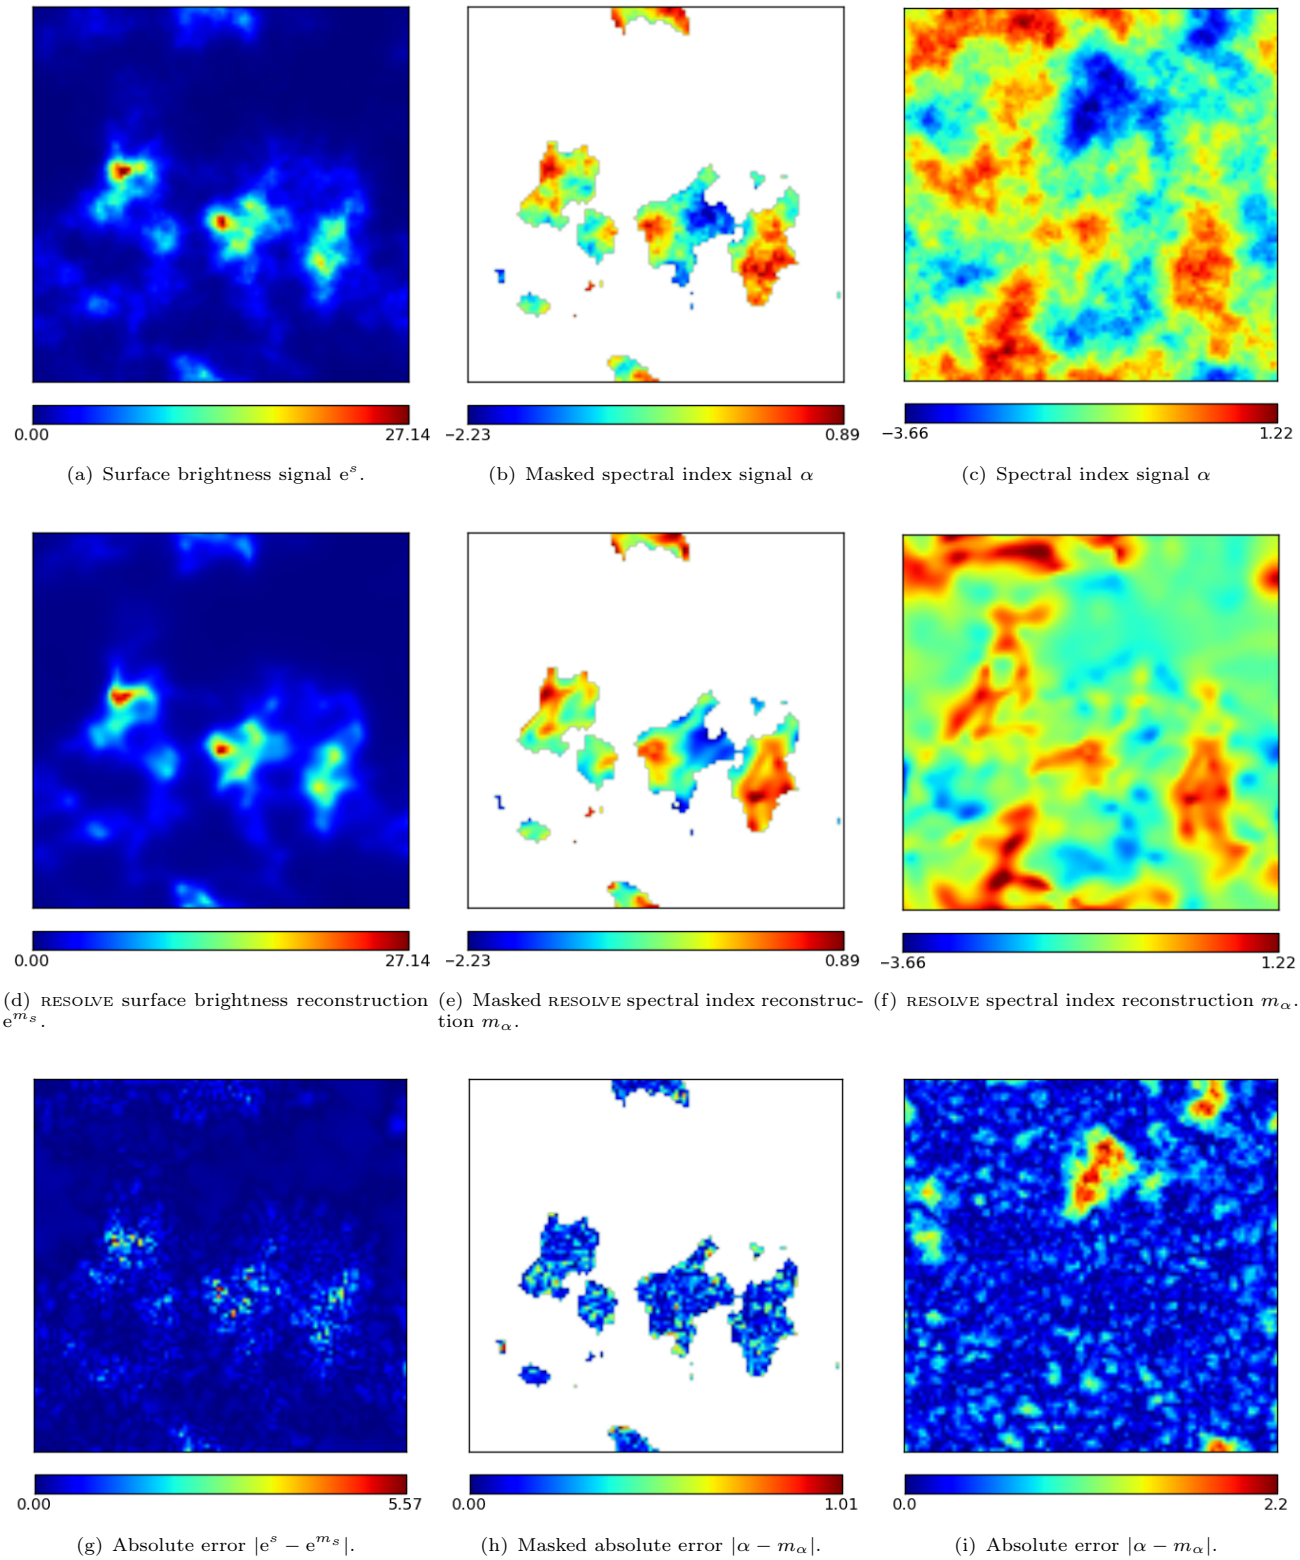

(a) Surface brightness signal e^s.
(b) Masked spectral index signal α
(c) Spectral index signal α
(d) RESOLVE surface brightness reconstruction e^{m_s}.
(e) Masked RESOLVE spectral index reconstruction m_α.
(f) RESOLVE spectral index reconstruction m_α.
(g) Absolute error $|e^s - e^{m_s}|$.
(h) Masked absolute error $|\alpha - m_\alpha|$.
(i) Absolute error $|\alpha - m_\alpha|$.

Fig. A.1. Multi-frequency reconstruction of the two signal fields e^s and α, observed with a sparse uv-coverage from a VLA-A-configuration (see Fig. 1). The images are 100^2 pixels large, the pixel size corresponds to roughly 0.1 arcsec. The brightness units are in Jy/px. The ridge-like structures in the difference maps simply stem from taking the absolute value and mark zero-crossings between positive and negative errors.

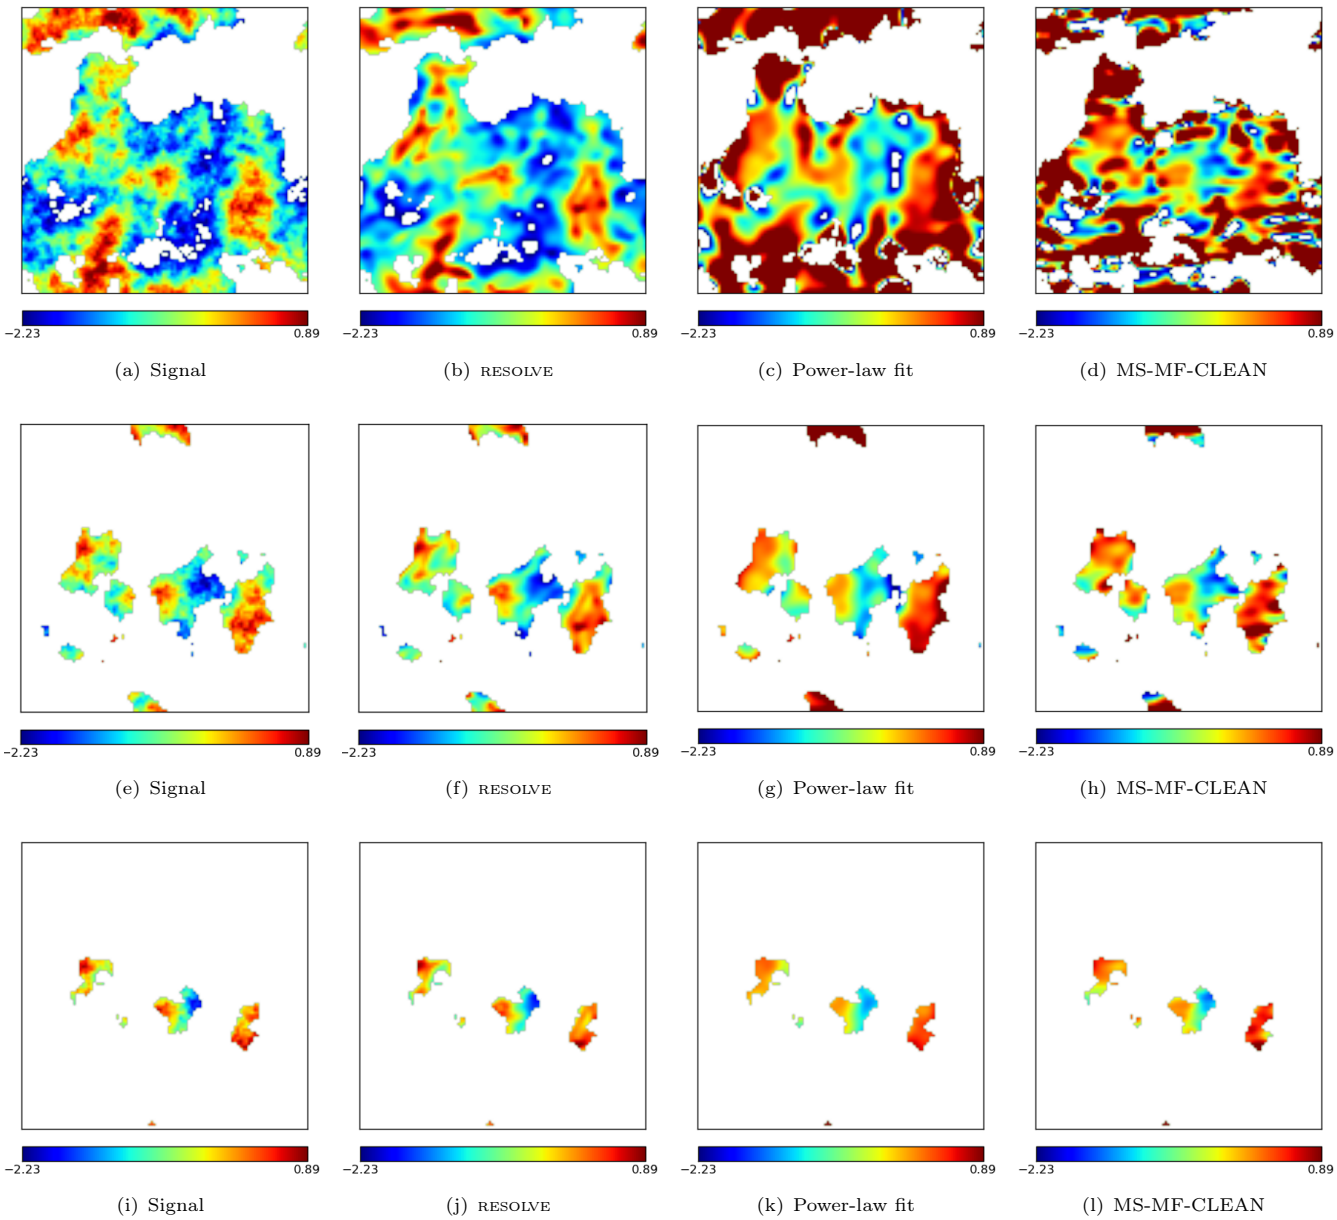

(a) Signal (b) RESOLVE (c) Power-law fit (d) MS-MF-CLEAN

(e) Signal (f) RESOLVE (g) Power-law fit (h) MS-MF-CLEAN

(i) Signal (j) RESOLVE (k) Power-law fit (l) MS-MF-CLEAN

Fig. A.2. Comparison of different methods for spectral index reconstruction with differently strong masks. The images are 100^2 pixels large, the pixel size corresponds to roughly 0.1 arcsec. *First column*: Signal. *Second column*: RESOLVE reconstruction. *Third column* Power-law fit. *Fourth column*: MS-MF-CLEAN.

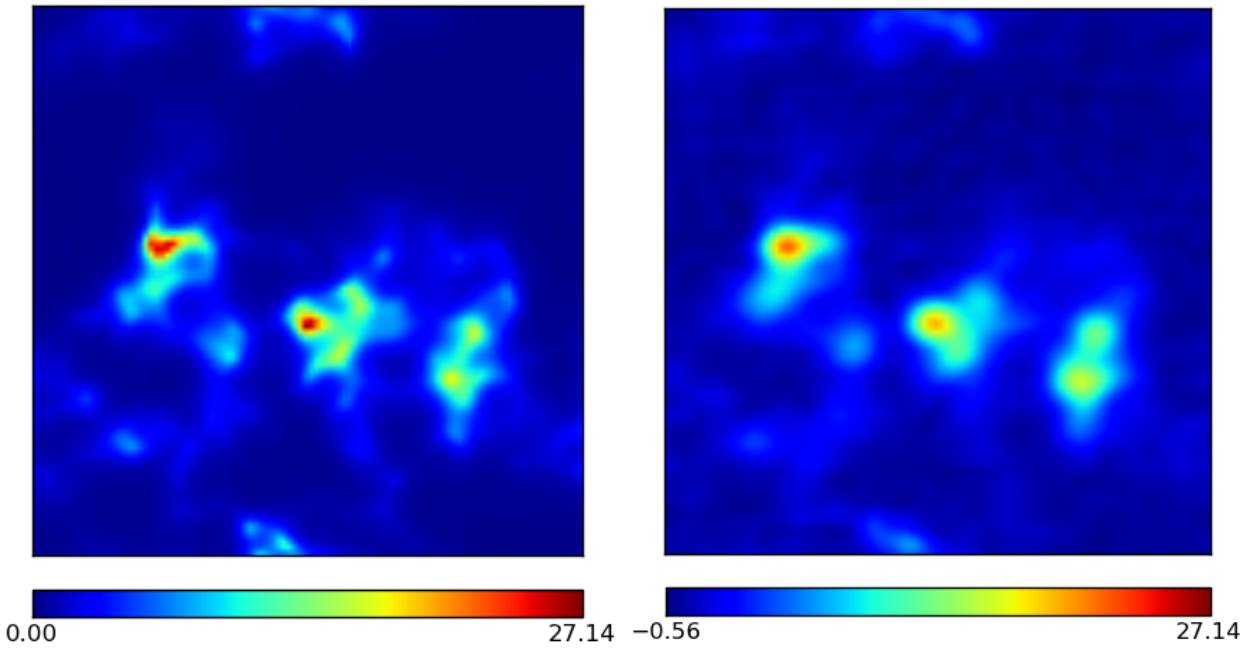

(a) RESOLVE brightness reconstruction e^{m_s}

(b) MS-MF CLEAN brightness reconstruction

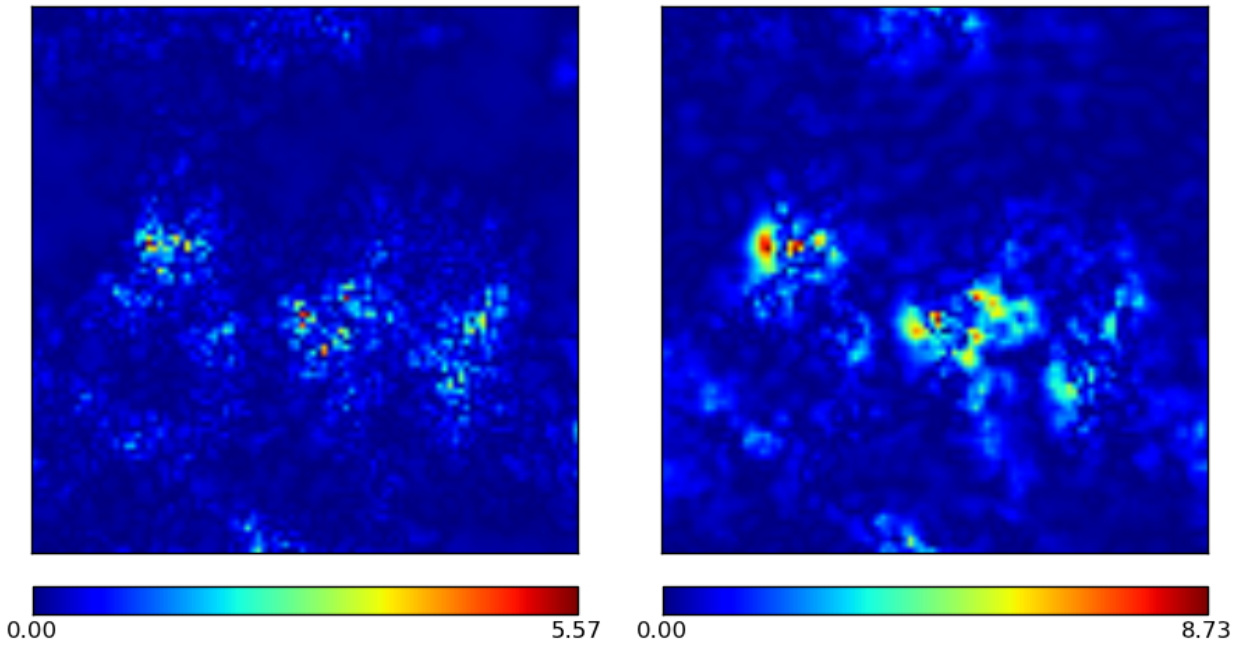

(c) Absolute error $|e^s - e^{m_s}|$

(d) Absolute error $|e^s - m_{\mathrm{clean}}|$

Fig. A.3. Comparison of multi-frequency RESOLVE and MS-MF-CLEAN surface brightness reconstructions. The images are 100^2 pixels large, the pixel size corresponds to roughly 0.2 arcsec. The brightness units are in Jy/px. The ridge-like structures in the difference maps simply stem from taking the absolute value and mark zero-crossings between positive and negative errors. *First row left*: RESOLVE reconstruction . *First row right*: MS-MF CLEAN reconstruction. *Second row left* Absolute per-pixel difference between the signal and the RESOLVE reconstruction *Second row right*: Absolute per-pixel difference between the signal and the MS-MF CLEAN reconstruction.